

In the Specification:

Please substitute the following paragraphs for the corresponding paragraphs beginning at the indicated location in the specification as originally filed.

(Page 9, line 12+)

~~Figure 3, comprising portions Figures 3a, 3b, and 3c, is a graphical interpretation~~ are graphical interpretations of the data used in the detection of novelty from the right visual field of an exemplary robot.

(Page 12, line 12+)

The robot 21 faces a demanding perceptual problem in determining what constitutes an obstacle 20 without being explicitly taught as both terrain with and without obstacles produce complex patterns of visual stimuli. In the present invention, an obstacle 20 becomes implicitly defined as any potentially destabilizing element of the environment. If the robot collides with the environment, it must refer back to the Sensorimotor Map 36 to determine what it saw previously and use that information (e.g. as seen in ~~portion Figure 3c of Figure 3~~) to adjust its control system not to make the same mistake again.

(Page 14, line 3+)

Referring now to ~~Figure 3 comprising portions Figures 3a, 3b, and 3c, an exemplary process of~~ a dynamic attention mechanism is shown that operates in a way such as to detect unexpected visual stimuli based on the state of all perceptual information and the locomotor controller (e.g. joint commands, tactile, disparity, and phase of gait information). ~~Figure 3 is accordingly divided into Figures 3a, 3b, and 3c represent~~ the three key layers of operation of the dynamic attention mechanism: raw data input ~~3a~~ (Figure 3a), prediction ~~3b~~ (Figure 3b), and novelty detection ~~3c~~ (Figure 3c). The process is demonstrated using data from the right side of the visual field only, but the procedure is identical for the left side as well.

(Page 14, line 16+)

In the raw data layer shown in ~~portion 3a of Figure 3~~ Figure 3a, the activation of the right vector cells, with eighteen (18) elements, versus phase of gait (described as θ in Figure 5), divided into twenty (20) discrete segments (each representing 1/20th of a gait cycle), for a total of three-hundred-and-sixty (360) cells 38. A gait cycle can be defined by assigning an arbitrary point as the

beginning of the gait cycle. The unfolding trajectory until the beginning of the next gait cycle (reaching that same arbitrary point of motion) can be parameterized by a single variable called phase. Cells 38 with lower numbers 38a, according to the graph, are closer to the robot whereas cells with higher numbers 38b are further away. The array of cells appears inclined consistent with a view of the surface from above at an oblique angle. Undulation in the phase direction corresponds to viewing height change during walking. Other more random variations thus represent perceived (Figure 3a) or predicted (Figure 3b) surface irregularities or possible obstacles.

(Page 15, line 3+)

In an exemplary prediction layer shown in ~~portion 3b of Figure 3~~ Figure 3b, the graphical representation of the predicted appearance of a surface is organized in cells 38 by disparity and phase as in Figure 3a ~~portion 3a of Figure 3~~. Each cell 38 receives information about an area of the surface from all sensors 31, 32, 33, (preferably encoded in a sparse code). The weight for each signal is determined by a learning rule (e.g. Widrow-Hoff LMS associative learning rule, etc.). The learning rule chosen is a supervised learning neural network learning rule although it may be possible to achieve the same results with an unsupervised learning rule as well. The primary function of the learning rule is to change the input weights of each cell such that it becomes a better predictor of sensor stimuli as time progresses. The learning rule reduces the weight from sensors with little predictive value and increases those with greater predictive value. The prediction is generated by a weighted average of all sensory and motor data (phase, motor signals (efference copies) tactile sensation, etc.) This adaptation is continuous through the 'life' of the robot 21 as the overall architecture is robust against loss of any sensor modality as all sensory information contributes to prediction of each other sensor.

(Page 15, line 30+)

As arranged in Figures 3a, 3b, and 3c ~~Figure 3~~, an exemplary novelty layer ~~3c (Figure 3c)~~, receives the difference between the raw data layer ~~3a (Figure 3a)~~ and the prediction layer ~~3b (Figure 3b)~~ weighted by a variable gain factor in order to determine an obstacle without being explicitly taught. The gain factor 42 for novelty detection varies due to a local feedback mechanism. The gain adjusts to maintain a low average activity at all times. If a certain cell has little predictive value, the cell's gain is reduced. If other cells predict the actual sensory input very

accurately, that cell's gain is increased, allowing finer discrimination. The output function of the novelty layer 41 represents a hard-limit threshold.

(Page 16, line 10+)

Thus, the dynamic attention mechanism, ~~shown in Figure 3~~ comprising procedural steps shown in Figures 3a, 3b, and 3c, allows the robot 21 to detect fine environmental features (e.g. an obstacle 20) of 1 cm in height or less whereas without the predictive component of this mechanism, the otherwise same device could not reliably detect obstacles less than 5 cm in height. As such, even small disparities between the actual/perceived (e.g. raw data layer 40) and predicted features (e.g. prediction layer 39), that correspond to just a fraction of a disparity value are recognized as novelty. Disparity has been defined by those well versed in the field of binocular vision and stereopsis as the side to side (horizontal) or up and down (vertical) "difference in the position of similar images in the two eyes...and can produce a compelling sensation of three-dimensionality." In this implementation, disparity values can easily vary +/-1 disparity value for a particular cell during walking and 3-4 disparity values between cells. Learning converges quite rapidly using this method such that good predictions and expectancy are obtained within one-hundred-and-twenty (120) seconds after initiation.

(Page 16, line 33+)

More particularly, Figure 7 illustrates the pseudo-cerebellum 30 in which the dynamic attention mechanism (~~Figure 3~~) functions. The pseudo-cerebellum 30 reacts to the information derived from the distal sensor(s) 33 to perform dynamic attention mechanism functions in each of the subregions 43 of the pseudo-cerebellum. Each subregion predicts sensory information based on both visually geometric stimuli including optic flow and other distal cues, as well as tactile stimuli and vestibular stimuli. The stimuli of each subregion 43 is in terms of distance (e.g. near stimuli to far stimuli). Within each subregion 43, prediction 39 is made in consideration of an efference copy 44 and other sensory input 31,32 using the formula

$$f((x \cdot w) - t)$$

where f is the neural output, x is a vector of inputs, sparsely coded, w is a vector of weights, and t is a threshold value. The function must have a non-linear form

and can be as simple as $f(x) = \max(0, x)$, a sigmoidal function or a $\tanh(x)$ function.

(Page 18, line 6+)

After a brief learning period, the robot 21 can accurately predict novelty 41 based on afferent responses. As illustrated in Figure 2, the information collected by the pseudo-cerebellum 30 using the dynamic attention mechanism (Figure 3) is then processed for association using a Sensory Motor Map 36, which can be modified later in the event an error occurs. The sensory data that is referenced as having resulted in that error will then be recognized and avoided to refrain from future repetition of that error. As a result of the robot's learning capabilities and expectancy, the robot 21 also learns to expect a smooth surface in front of it when trained on a smooth surface, and without being explicitly told about smooth surfaces (or a rough surface when trained on a rough surface, without being explicitly taught about rough surfaces).

(Page 20, line 14+)

The robot 21 can learn to adjust stride length based on an activated novelty cell (e.g. something other than predicted) triggers an eligibility trace. An eligibility trace is a short-term memory delay signal which allows association between future and current events; if the robot's foot collides with the environment, a training signal 45 (representing error) is sent to a sensorimotor mapping mechanism 36 from the novelty cells (shown in ~~portion 3c of Figure 3~~ Figure 3c) to a variable that adjusts stride length in the CPG 29.